

Oscillations in the Evolution of Reciprocity

MARTIN NOWAK

*Institut für Theoretische Chemie der Universität Wien, Währingerstr. 17,
A-1090 Wien, Austria*

AND

KARL SIGMUND

*Institut für Mathematik der Universität Wien, Strudlhofg. 4, A-1090 Wien,
Austria, and IIASA, Laxenburg, Austria*

(Received 20 August 1988)

A game-theoretical analysis of the Iterated Prisoner's Dilemma shows that the evolution of ensembles of stochastic strategies displays a dynamics of high complexity and unpredictability.

An increasing amount of attention has been paid to the Iterated Prisoner's Dilemma as a game theoretical paradigm for the evolution of co-operation based on reciprocity. In repeated encounters, two individuals are faced with the options to co-operate or to defect (*C* or *D*). Joint co-operation leads to a payoff *R* (reward) which is higher than the payoff *P* (punishment) for mutual defection. But if one player "cheats" by defecting while the other co-operates, then his payoff *T* (temptation) is larger than *R*, while the co-operator's payoff *S* (sucker) is smaller than *P*. In addition to $T > R > P > S$ one usually assumes $R > \frac{1}{2}(S + T)$ in order to simplify the analysis by making alternations of co-operation and defection less rewarding than all-out co-operation.

If the game consists of a single encounter (or of a fixed number of encounters known to both players) then the best choice is to defect. But if the length of the game is unknown, as for example if there is a fixed probability *w* for a further encounter, then the players may find that it is in their interest to co-operate. In Axelrod's well known computer tournaments (Axelrod, 1984) nice, i.e. co-operative, strategies did very well, and the simplest one finished at the head of the class. This was Rapaport's Tit-For-Tat (*TFT*), a strategy which starts with a co-operative move and then does whatever the opponent did on his previous move.

The assessment in Axelrod's contests was originally established by round-robin tournaments. For applications to evolutionary biology, an "ecological approach" was proposed (Axelrod & Hamilton, 1981), where each strategy participates to the next generation in proportion to its present success. Thus good strategies spread in the population and eliminate weaker ones: but the success of a strategy depends not only on its own merits but also on the frequencies of its competitors. There are several ways to model such an evolution, essentially leading to similar results. We

shall use here the continuous game dynamics of Taylor & Jonker (1979) (see Hofbauer & Sigmund (1988) for a survey).

In spite of its success, *TFT* is not evolutionarily stable in the sense of Maynard Smith (1982). It was shown (Axelrod, 1984) that for sufficiently high w , *TFT* cannot be invaded by All Defect (*ALLD*), but there are many strategies (*ALLC*, for example) which do as well as *TFT* against each other and against *TFT*. They can invade, not by selection pressure but by genetic drift. Once they are common in the population, less nice strategies can spread, because they have to fear less retaliation than against *TFT* alone (Selten & Hammerstein, 1984).

Another argument (Boyd & Lorberbaum, 1987) used is that if a strategy like Suspicious Tit-For-Tat (*STFT*, which defects on the first move and then does like *TFT*) is maintained by mutation pressure in a *TFT*-population, then Tit-For-Two-Tats (*TFTT*, which defects only after two consecutive *D*'s by the opponent) can invade. Indeed, it does like *TFT* against itself and against *TFT*, but better against *STFT*. This is not properly an argument against it being evolutionarily stable (using mutations to maintain *STFT* is not quite in the rule-book), but it further weakens the claim that *TFT* cannot be invaded. May (1987) stresses rightly that the success of strategies should realistically be evaluated in the presence of "representative ensembles" of strategies, and not just against a single mutant.

May also points out that more account should be taken of stochasticity. Actual biological situations are fraught with errors and uncertainties. The answer to the opponent's last move (which may be misperceived in the first place) is never an all-or-nothing, but only an increase or decrease of the propensity to co-operate. This emerges quite clearly from Milinski's neat experiments on sticklebacks (Milinski, 1987) or Lombardo's data on swallows (Lombardo, 1985). Even the smallest random fluctuation can have drastic effects, for example in encounters between *TFT*-players, which can get locked, by a single mistake, into a series of alternating defections which can only be broken by another error.

This suggests considering stochastic strategies given by three parameters (y, p, q) , where y is the probability to co-operate on the first move, and p and q the conditional probabilities to co-operate, given that the opponents last move was a *C* or a *D*. This class contains for example *TFT*(1, 1, 0), *STFT*(0, 1, 0), *ALLC*(1, 1, 1) and *ALLD*(0, 0, 0) as extremal representatives. It does not contain *TFTT*, or strategies which also depend on one's own last move, or the strategies determined by the history of the last three moves which were used in Axelrod's simulations by genetic algorithms (Axelrod, 1987).

In spite of this, strategies of type (y, p, q) already display a remarkable variety of evolutionary dynamics. (For our examples, we use Axelrod's values $R = 3$, $T = 5$, $P = 1$ and $S = 0$ for the payoffs, and $w = 0.9$ as discount parameter. Other choices display similar behaviours.)

If only two strategies are competing, one can find dominance (e.g. $S_1 = ALLD$ always outcompetes $S_2 = ALLC$), bistability ($S_1 = ALLD$ and the stochastic *TFT* $S_3 = (y, 1, 0)$, with $0 < y < 1$, do never co-exist, but which one wins depends on the initial frequencies) and stable polymorphism ($S_2 = ALLC$ and $S_3 = (y, 1, 0)$ settle down to an equilibrium). (We mention that Feldman & Thomas (1987) have also

found polymorphism if the probability w for continuing the game is part of the strategy.)

For three competing strategies, it may depend on the initial condition whether a polymorphic state gets established or not. An example is obtained by the three strategies $S1$, $S2$ and $S3$ above: most initial conditions lead to an equilibrium of all three strategies which is, however, not evolutionarily stable (see Fig. 1). For $w = 1$, one finds a "stone-scissors-paper" cycle: $S1$ is dominated by $S3$, $S3$ by $S2$ and $S2$ in its turn by $S1$. The state space is filled by neutral oscillations in this case. If $w < 1$ but $y = 1$ then most initial conditions yield a mixture between $ALLC$ and TFT , with $ALLD$ eliminated. One finds "stone-scissors-paper" cycles for $w < 1$, too: for example

(a) $S1(0.40, 0.75, 0.75)$, $S2(0.40, 0.75, 0.25)$, $S3(0.40, 0.95, 0.25)$

or

(b) $S1(0.75, 0.75, 0.75)$, $S2(0.75, 0.75, 0.25)$, $S3(0.54, 0.95, 0.30)$.




FIG. 1. Phase diagram for the frequencies x_1, x_2, x_3 of the strategies $S1(0, 0, 0)$, $S2(1, 1, 1)$, $S3(0.9, 1, 0)$ on the simplex. Two trajectories are shown: one converges to $S1$, the other to a stable polymorphism.

If all three strategies are initially present, then in case (a) the frequencies of S_1 , S_2 and S_3 converge in an oscillatory way to a stable polymorphism, while in case (b) they oscillate with increasing amplitude and exponentially decreasing frequency. This case is sketched in Fig. 2: for some time, S_1 seems to predominate, until it is suddenly replaced by S_2 , which seems to prevail for some longer time, until it is brusquely displaced by S_3 , which in turn takes over for a still longer time, until it is shouldered away, apparently without exterior cue, by S_1 again etc., in a cyclic series of fits and starts. Analytically, such a "heteroclinic cycle" whose relevance to ecological models was first stressed by May & Leonard (1975), does never settle down. The preponderance of the topmost species becomes more and more extreme, its time of prevalence grows exponentially, but the sudden upheavals always take the same very short time. In practice, a random fluctuation, or in numerical experiments, a round-off term, will eventually wipe out one of the strategies, which means extinction for the next one and fixation for the third, but it is impossible to predict which one will end up as the winner.

With four competing strategies, one can find oscillations which damp down to some equilibrium, or which "explode" in the aforementioned way, or which settle down to some predetermined amplitude and period. Such a "limit cycle" is found for

$$(c) \quad S_1(0.75, 0.75, 0.75), \quad S_2(0.75, 0.75, 0.25), \quad S_3(0.40, 1.00, 0.30), \\ S_4(0.70, 1.00, 0.00)$$

(see Fig. 3).

In examples (b) and (c), there is no evolutionarily stable equilibrium. The (unique) Nash equilibrium is unstable. We cannot estimate the probability for such situations, but it seems to be fairly high. It is difficult to predict which set of strategies leads to complex dynamics, but it seems that at least one strategy should be a neighbour of *TFT*.

In the face of these intricacies, it seems clear that May's (1987) exhortation to "take more account of intrinsic stochasticities and of evolutionary stability against representative ensembles of mutant strategies" is a fairly tall order. One can approach it by numerical simulations, starting with an arbitrary distribution of strategies and introducing from time to time a mutant close to the prevailing ensemble. If one considers only strategies differing in a single parameter, the situation is reasonably clear:

- (i) variation of y leads, depending on p and q , either to a pure state with $y = 0$ or $y = 1$ or to a rich mixture of strategies with a predetermined average value for its initial readiness to co-operate;
- (ii) variation of p leads to an extremal value 0 or 1, depending on the initial state of the population;
- (iii) variation of q leads to a monomorphic population with a predetermined q -value depending on p and y .

If one admits variation in all three parameters, however, then the result is considerably less predictable. It depends obviously a lot on the initial conditions and on the history of mutational events. The outcome often is a population near





FIG. 2. A heteroclinic cycle for $S1(0.75, 0.75, 0.75)$, $S2(0.75, 0.75, 0.25)$, $S3(0.54, 0.95, 0.30)$: (a) a trajectory on the simplex; (b) the oscillation of x_2 in time.




FIG. 3. A limit cycle for $S1(0.75, 0.75, 0.75)$, $S2(0.75, 0.75, 0.25)$, $S3(0.40, 1.00, 0.30)$, $S4(0.70, 1.00, 0.00)$, shown in projection on the x_1 - x_4 -plane.

the *TFT*-value, at least in its time average. This underlines the robustness of *TFT*. A statistical analysis is still lacking, however, and an intuitive understanding is made difficult by the prevalence of oscillations.

This work was partly supported by the Austrian Forschungsförderungsfonds P6866.

REFERENCES

- AXELROD, R. (1984). *The Evolution of Cooperation*. New York: Basic Books.
- AXELROD, R. (1987). The Evolution of Strategies in the Iterated Prisoner's Dilemma. In: *Genetic Algorithms and Simulated Annealing* (Davis, D., ed.) pp. 32-42. London: Pitman.
- AXELROD, R. & HAMILTON, W. D. (1981). *The Evolution of Cooperation*. *Science, N.Y.* **211**, 1390-1396.
- BOYD, R. & LORBERBAUM, J. P. (1987). No pure strategy is evolutionarily stable in the Repeated Prisoner's Dilemma. *Nature, Lond.* **327**, 58-59.
- FELDMAN, M. & THOMAS, E. (1987). Behavior-dependent contexts for repeated plays of the Prisoner's Dilemma II: Dynamical aspects of the evolution of cooperation. *J. theor. Biol.* **128**, 297-315.
- HOFBAUER, J. & SIGMUND, K. (1988). *The Theory of Evolution and Dynamical Systems*. Cambridge: Cambridge University Press.
- LOMBARDO, M. P. (1985). Mutual Restraint in Tree Swallows: A Test for the TIT FOR TAT Model of Reciprocity. *Science, Lond.* **227**, 1363-1365.
- MAY, R. M. (1987). More evolution of cooperation. *Nature, N.Y.* **327**, 15-17.
- MAY, R. M. & LEONARD, W. (1975). Nonlinear aspects of competition between three species. *SIAM J. Appl. Math.* **29**, 243-252.
- MAYNARD-SMITH, J. (1982). *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.
- MILINSKI, M. (1987). Tit For Tat in sticklebacks and the evolution of cooperation. *Nature, Lond.* **325**, 434-435.
- SELTEN, R. & HAMMERSTEIN, P. (1984). Gaps in Harley's argument on evolutionarily stable learning rules and in the logic of *TFT*. *Beh. Brain Sci.* **7**, 115-116.
- TAYLOR, P. & JONKER, L. (1979). Evolutionarily stable strategies and game dynamics. *Math. Biosci.* **40**, 145-156.